

SYSTEM FOR EXTRACTION OF REPRESENTATIVE DATA FOR
TRAINING OF ADAPTIVE PROCESS MONITORING EQUIPMENT

5

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of priority under 35 U.S.C. § 119(e) to U.S. Provisional application serial no. 60/187,959 filed March 9, 2000.

10

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to equipment and process monitoring, and more particularly to monitoring systems instrumented with sensors that measure correlated phenomena. The present invention further relates to modeling instrumented, real-time processes using the aggregate sensor information to ascertain information about the state of the process, and a method of training an empirical model used therein.

20 2. Description of the Related Art

Conventional methods are known for monitoring equipment or processes – generically “systems” – using sensors to measure operational parameters of the system. The data values from sensors can be observed directly to understand how the system is functioning. Alternatively, for unattended operation, it is known to compare sensor data values against stored or predetermined thresholds in an automated fashion, and generate an exception condition or alarm requiring human intervention only when a sensor datum value exceeds a corresponding threshold.

A number of problems exist with monitoring systems using thresholds. One problem is the difficulty of selecting a threshold for a dynamic parameter that avoids a burdensome number of false alarms, yet catches real alarms and provides sufficient warning to take corrective action when a system parameter – as measured by a sensor – moves outside of acceptable operation. Another problem is posed by sensor failure, which may result in spurious parameter values. It may not be clear from a sensor

data value that the sensor has failed. Such a failure can entirely undermine monitoring of the subject system.

In systems with a plurality of sensors measuring correlated phenomena in the system, it is known to use certain methods to consider all 5 sensors in aggregate to overcome some of these problems. By observing the behavior of all the sensor data values in aggregate, it can be possible to dramatically improve monitoring without suffering unduly from false and missed alarms. Also, knowledge of how all the correlated parameters behave in unison can help determine that a sensor has failed, when isolated 10 monitoring of data from that sensor would not in and of itself indicate the sensor failure.

Known methods for viewing aggregate sensor data typically employ a modeling function that embodies prior knowledge of the system. One such technique known as a "first-principles" model requires a well- 15 defined mathematical description of the dynamics of the system selecting system snapshots taken at minimum and maximum system parameter excursions. The mathematical model is used as a reference against which current aggregate sensor data can be compared to view nascent problems or sensor failures. However, this technique is particularly vulnerable to even 20 the slightest structural change in the observed system and may not provide sufficient system characterization in operating regions where system parameters vary most dynamically. The mathematical model of the system is often very costly to obtain, and in many cases, may not be reasonably possible at all.

Another class of techniques involves empirically modeling the 25 system as a "black box", without discerning any specific mechanics within the system. System modeling using such techniques can be easier and more resilient in the face of structural system changes. Modeling in these techniques typically involves providing some historic sensor data corresponding to desired or normal system operation, which is then used to 30 "train" the model.

One particular technique is described in U.S. Patent No. 5,987,399, the teachings of which are incorporated herein by reference. As taught therein, sensor data is gathered from a plurality of sensors measuring

correlated parameters of a system in a desired operating state. This data is used to derive an empirical model comprising certain acceptable historical system states. Real-time sensor data from the system is provided to a modeling engine embodying the empirical model, which computes a 5 measure of the similarity of the real-time state to all prior known acceptable states in the model. From that measure of similarity, an estimate is generated for expected sensor data values. The real-time sensor data and the estimated inspected sensor data are compared, and if there is a discrepancy, corrective action can be taken.

10 Other empirical model-based monitoring systems are disclosed in US Patent No. 4,937,763 to Mott, wherein learned observations are employed in a system state analyzer, and U.S. Patent No. 5,764,509 to Gross et al., the teachings of which are hereby incorporated by reference. Selection of the appropriate historical sensor data for generating any of these 15 empirical models is a serious hurdle. The models variously rely on the historic data accurately representing the "normal" conditions of the process or machine being monitored. Therefore, one must ensure that the data collected as historic data corresponds to an acceptable state of operation, and not one in which a latent fault was present in the process or machine. A 20 larger problem is then to ensure that the historic data is sufficiently representative of the expected ranges of operation, so that the empirical model does not generate alarms for states of operation it has no history for, but which are otherwise acceptable states for the process or machine. It is critical to the success of the empirical model for monitoring that the 25 collected sensor data be properly distilled or condensed to a trained set of data that adequately represents the knowledge of the normal states of operation of the process or machine being monitored. An additional problem is that, since empirical modeling methods can be computationally demanding, it is often preferable to restrict the historic data on which they 30 are built or trained to a minimum, in order to reduce training time and required computing power. Finally, some empirical models are actually adversely affected by too much training data: They tend to find every current state of the monitored process or machine acceptable, because something close enough to it can be found in the historic data. Therefore, a

successful selection of representative "training set" data must not result in an "overtrained" model.

In US Patent No. 5,764,509 to Gross et al., a training method for selecting observations of time-correlated sensor data called Min-Max is presented. According to this way of training a model, the collected normal sensor data is condensed or distilled down to a "training set" by selecting those observations (or "snapshots") that contain a global maximum or minimum for a sensor with respect to all values taken on by that sensor across the entire collected sensor data. Thus, as a maximum the number of observations that are include in the training set that results from the training is twice the number of sensors being modeled. While this method assures the inclusion of extrema for all sensors in the model, it may be desirable to enhance the model with inclusion of other snapshots with intermediate values.

Therefore, when selecting vector-arranged snapshot data for inclusion in a training set for deriving an empirical model, there is a need for selecting an optimized training set that best characterizes the dynamics of the underlying machine or process. There is a further need for a method for selecting historic data that minimizes the size of the training set. Finally, there is a need for training methods that are computationally efficient and fast. This invention achieves these benefits by automating selection in a way that maximizes the data membership from regions of great dynamics, while keeping the overall training set size manageable.

25 SUMMARY OF THE INVENTION

The invention provides several benefits for building a representative training set from a larger data collection for empirical modeling of a process or machine. In generating an empirical model for monitoring, controlling or optimizing the operation of a process or machine (industrial or otherwise), the invention provides a novel and improved means for identifying and selecting a training set from a large volume of historic data of the process or machine. Historic data is collected as snapshots of time-correlated sensor data from the process or machine. The sensors can be of any type, measuring any kind of physical or derived

parameter. The collected data can be provided in time sequence or out of sequence without affecting the results produced by the invention.

Briefly summarized, the present invention selects appropriate modeling data from a collected data set to characterize a modeled process.

5 Typically, the data is correlated sensor data, representing parametric conditions at a multitude of snapshots of a system, machine or process parameters. More observations are selected for inclusion in the model for ranges of sensor values where that data exhibits greater dynamics, while including less data from stable ranges of little change. The system of the
10 present invention can be a computer running a program in software and the computer may simply be a microprocessor.

According to the present invention, the distribution of chosen vectors may be varied to suit the specific needs within a training set. For example, selected training set vectors can be evenly spaced along the value 15 range of a specific variable on the y axis, or chosen as population of controlled entities by using a distribution along the x axis. Several nonlinear variations of the invention may be applied to either axis, including a Gaussian distribution, grouping vectors into narrow ranges coupled with random selection from the entire vector population, and randomly selecting 20 n vectors from each range, where n is a function of vector population. Thus, a set of vectors may be chosen to fully represent the range of each variable, providing full range modeling capability, while including a number of closely spaced observations (represented by the selected vectors) centered about the mean of the data. Conversely, the resolution can be enhanced 25 selecting only a few samples at the mean where variation is expected to be fairly minor to allow greater precision in specification of values at the edges of the data set. Thus, very high model fidelity is achieved for normal operation with minimal growth of the training matrix or the G matrix derived therefrom. In addition, it is possible to use different criteria for each 30 variable in the observation vector, giving the user great customization capability.

Thus, it is a purpose of the invention to automate sensor snapshot selection to maximize the data membership from regions of great dynamics, while keeping the overall training set size manageable;

It is another purpose of the invention to selectively highlight certain variables or parameter sensors thereby weighing some more than others and so, emphasizing certain data regions more in the training set;

5 It is yet another purpose of the invention to provide the flexibility to work with many different data types and many different relationships between the variables within a data collection.

10 Advantageously, the present invention is particularly valuable when it is important to tune the operation of a similarity operator to meet requirements which include memory footprint constraints and model performance is important.

The described embodiment is a system that employs a set of user-defined parameters to select a training set from a data collection. The invention may be enabled in performing the selection process in concert with a software or hardware based data storage system.

15 Preferably, each parameter variable in a data collection is treated equally with every other parameter or variable and the same number of observation points are selected from each at equally spaced intervals of magnitude.

20 BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as the preferred mode of use, further objectives and advantages thereof, is best understood by reference to the following detailed description of the 25 embodiment in conjunction with the accompanying drawings, wherein:

FIG. 1 shows an empirical model-based monitoring system for which a training set may be derived according to the preferred embodiment of the present invention;

30 FIG. 2 shows an example wherein two independent vectors are checked for similarity using the prior art BART technique;

FIG. 3 is a flowchart of a method of generating and employing an empirical model for process or machine monitoring;

FIG. 4 shows a prior art method for selecting training set vectors for distilling the collected sensor data to create a representative training data set;

5 FIG. 5 graphically depicts a sequence of values as a bar chart ordered with increasing parameter height;

10 FIG. 6A shows a block diagram of an exemplary laboratory workbench arrangement for gathering process or machine behavior data for distillation;

15 FIG. 6B shows an example of an exemplary monitoring system with an on-board processor and a training set selected according to a preferred embodiment of the present invention;

20 FIG. 7 shows an example wherein a monitored process is shown to be instrumented with sensors having output leads and a training set derived according to the present invention is used in monitoring the system;

25 FIG. 8 shows a cumulative distribution function for the same data as in FIG. 5;

30 FIG. 9 shows a flow diagram for selecting a training set as in the example of FIG. 5.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention is a method system and program product for distilling a training set from a collection of data. According to the invention, for each sensor or parameter provided in the model, all 25 collected snapshots are arranged in a sequence that orders the sensor of interest according to magnitude. The ordered snapshots are then chosen for inclusion in the final training set used as the basis for the empirical model, by segmenting the magnitude axis into equal-spaced segments, and identifying one snapshot for each segment. This is performed for each 30 sensor. More specifically, the magnitude axis is divided into equal segments, and snapshot with the magnitude for the sensor of interest that is closest to any segment divider value is included in the training set in its entirety. The training set selection can be done on processed or unprocessed data. The data is analyzed and the training set is selected by dividing the

data, uniformly or non-uniformly, into as many discrete bins as would yield the desired size for the training set. Various nonlinear options may be selectively included for focusing the behavior of the resulting model to suit specific application needs.

5 The invention is beneficially understood in the context of the empirical model-based monitoring system for which it can provide a training set. Turning to FIG. 1, such a monitoring system is shown to comprise a data acquisition module 102, an information processor 104, a memory 106 and an output module 108, which can be coupled to other 10 software, to a display, to an alarm system, or any other system that can utilize the results, as may be known in the art. The processor 104 generally may include a similarity engine 110, an estimated state generator 112 and a deviation engine 114.

15 Memory 106 stores a plurality of selected time-correlated snapshots of sensor values characterizing normal, optimal, desirable or acceptable operation of a monitored process or machine. This plurality of snapshots, distilled according to a selected "training" method as described herein, comprises an empirical model of the process or machine being monitored. In operation, the inventive monitoring system 100 samples 20 current snapshots of sensor data via acquisition module 102. For a given set of time-correlated sensor data from the monitored process or machine running in real-time, the estimates for the sensors can be generated by the Estimated State Generator 112 according to:

$$\bar{y}_{estimated} = \bar{D} \bullet \bar{W}$$

25 where D is a matrix comprised of the plurality of snapshots in memory 106 selected according to training, and W is a contribution weighting vector determined by Similarity Engine 110 and Estimated State Generator 112 using a similarity operator such as the inventive class of similarity operators of the present invention. The multiplication operation is the standard 30 matrix/vector multiplication operator. W has as many elements as there are snapshots in D, and is determined by:

$$\vec{W} = \frac{\hat{W}}{\left(\sum_{j=1}^N \hat{W}(j) \right)}$$

$$\underline{\hat{W}} = \left(\vec{D}^T \otimes \vec{D} \right)^{-1} \bullet \left(\vec{D}^T \otimes \vec{y}_{in} \right)$$

where the T superscript denotes transpose of the matrix, and Y(in) is the current snapshot of actual, real-time sensor data. The symbol “ \otimes ” in the 5 equation above represents the improved similarity operator of the present invention. Y(in) is the real-time or actual sensor values from the underlying system, and therefore it is a vector snapshot.

The similarity operation typically returns a scalar value 10 between 0 and 1 for each comparison of one vector or matrix row/column to another vector. It represents a numeric quantification of the overall 15 similarity of two system states represented by two snapshots of the same sensors. A similarity value closer to 1 indicates sameness, whereas a similarity value closer to 0 typically indicates difference. A variety of techniques that implement a similarity operator are known, such as BART in US Patent No. 5,987,399, as well as that described in US Patent No. 5,764,509.

Deviation detection engine 114 receives both the actual current 20 snapshot of sensor values and the set of sensor value estimates from the estimated state generator 114, and compares the two. A variety of tests can be used, including the sequential probability ratio test, or a CUSUM test, 25 both of which are known in the art.

FIG. 2 illustrates the BART technique as described in U.S. Patent No. 5,987,399, the teachings of which are incorporated herein by reference, wherein triangle 200 is formed having a linear base 202 bounded in training data by the range for a given sensor, the range min and max 25 forming vertices 204, 206 at opposite ends of the base. The triangle 200 was formed as a right triangle, and the location of the right angle was located a height h above the median of the range data along the base 202. The height h was required to be chosen so that the apex angle is a right angle. To

5 perform a similarity operation on two snapshots, then, corresponding elements of the snapshots are individually compared for similarity using this triangle 200. For a given sensor, a first sensor value X_0 from the first snapshot and a second sensor value X_1 from the second snapshot are plotted along the base 202 according to where they fall between the minimum MIN and maximum MAX values for that sensor across the entire training set.

10 This generates an angle theta (θ) that is compared to angle omega (Ω). A relatively small θ indicates high similarity of the sensor values, whereas a large θ relative to Ω indicates low similarity. The elemental similarity for the pair of sensor values is averaged with analogously derived elemental similarities for other sensor value pairs in the two snapshots to provide an overall vector-to-vector similarity score.

15 The operation of the monitoring system can be better understood with examination of FIG. 3, which is a flowchart describing a method of generating and employing an empirical model for process or machine monitoring. In step 310, data is collected from the instrumented process or machine, while it is operating through all of its acceptable dynamic ranges. The data takes the form of snapshots of time-correlated data, which can be treated as one-dimensional arrays or vectors. This historic data may comprise a huge volume of snapshots. It is distilled in step 320 according to the invention to a subset of snapshots (a "training set") sufficiently representative of the expected dynamic ranges of operation of the process or machine. The training set may comprise a mere fraction of the original historic data set. In step 330, the training set is used to build an empirical model using one of a variety of available empirical modeling techniques. Once the empirical model is built, the data representing it can be loaded in step 340 to a real-time monitoring system, typically a computer platform adapted to received real-time sensor data from the process or machine. Real-time monitoring begins in step 350 with capturing a current snapshot of the same sensors or parameters that were used in the training set data. In step 360, the empirical model operates on the current snapshot to generate a set of estimates for the sensors or parameters. This is the estimation by the model of what values these sensors should have based on what is being reported. These estimates are compared to the actual current

sensor data in step 370. Alerts are generated as indicated and other post-processing for display or control purposes is carried on in step 380.

Monitoring continues with step 350, where the next current snapshot is captured.

5 The present invention may be used with technology like BART and that described in U.S. Patents No. 5,764,509 (MSET), the teachings of which are incorporated by reference. Such systems as MSET and BART employ empirical modeling in conjunction with a statistical hypothesis test to yield excellent sensitivity to incipient changes in the operational state of a
10 monitored process or machine. The statistical hypothesis test can be the Sequential Probability Ratio Test (SPRT), which accepts two inputs to determine if they are the same or different with statistical confidence. Real-time sensor data, or a source of actual parametric data, provides one input to SPRT. The empirical model provides another input, by generating an
15 estimate from the real-time sensor values or actual parametric data values.

 The empirical model is achieved by selecting past history data that reflects desired states of operation for a process or machine that is monitored with sensors and SPRT. The empirical model in a sense has “learned” the known operational states of the process or machine from this
20 historic data.

 Typically, when creating or “training” an empirical model as described in the aforementioned patents, a large amount of available data from the process or machine must be distilled to a computationally manageable “training” set. The training set must still be sufficiently
25 representative of the full dynamic ranges of the process or machine that the empirical model can render reasonable estimates in real-time for sensor values. Known methods for doing this include the aforementioned “Min-Max” procedure described in U.S. Patent No. 5,764,509. According to this procedure, data is selected that includes all minima and maxima for each
30 measured parameter included in vectors in the available data.

 FIG. 4 graphically depicts the MIN-MAX selection method of the prior art for distilling the collected sensor data to create a representative training data set. In this example, only five sensor signals 402, 404, 406, 408 and 410 are shown for a system, process or machine. On the abscissa axis

412 is the sample number or time stamp of the collected sensor data, where
the data is digitally sampled and the sensor data is temporally correlated.
The ordinate axis 414 represents the magnitude of each sensor reading at the
particular sample or snapshot. Each snapshot represents a vector of five
5 elements, each element corresponding to a reading for a sensor in that
snapshot. Of all the sensor data collected (in all of the snapshots), according
to the prior art training method, only those five-element snapshots are
included in the representative training set that contain either a global
minimum or a global maximum value for any given sensor. Therefore, for
10 sensor signal 402 the global maximum 416 qualifies the five sensor values
that intersect with line 418 including global maximum 416, for inclusion in
the representative training set as a five element vector. Similarly, for sensor
signal 402 the global minimum 420 qualifies the five sensor values that
intersect with line 422 for inclusion in the vector training set. So, collections
15 of such snapshots represent states the system has taken on and that are
expected to reoccur. The pre-collected vectors are further filtered to produce
a “training” subset that reflects all min/max states that the system is known
to take on while operating “normally” or “acceptably” or “preferably”. This
training set forms a matrix, having as many rows as there are sensors of
20 interest, and as many columns (snapshots) as necessary to capture all the
minimum and maximum states without redundancy.

The training matrix of observed acceptable states is referred to
as “D” and allows a computation of an expected snapshot given the real-
time actual snapshot of the underlying system in operation. The present
25 invention can be applied to this acceptable state matrix D, selectively
augmenting it with snapshots from system operating regions of particular
interest, e.g., dynamic regions of parameter sensitivity or at parametric
extremes. Thus, additional snapshots may be included in the acceptable
state matrix D, thereby, allowing determination of a much more refined and
30 precise expected snapshot.

These prior methods of automating data selection from
available historic data are adequate, but still have shortcomings. Min-Max
typically selects an insufficiently descriptive training set for optimal
modeling, for all but the least dynamic of systems. Vector Ordering

produces a training set much more representative of the dynamics of the system, but not in a way that is specific to particularly dynamic sensors.

The present invention overcomes these shortcomings, to produce a training set from historic available data that is fairly representative of the dynamics of the monitored process or machine, yet computationally manageable. Parameter data are gathered from signal sensors monitoring a system such as a machine, process or living system. The number of sensors used is not a limiting factor, generally, other than with respect to computational overhead. The present invention is highly scalable. The sensors should capture component parameters of at least some of the primary "drivers" of the underlying system. Furthermore, all sensor inputs to the system are best interrelated in some fashion (non-linearly or linearly).

As used herein, the word "sensor" is not intended to be restrictive but to indicate the source of system, process or machine parameter data and is intended to be taken to mean, generally, any kind of collected data taken and collected by any means. Thus, sensor data may include, for example, a digitally converted value from a real-time sensor, a sensor datum stored in a computer file, or some other type of data that may measure parameters of a system or process without actually requiring a physical sensor, e.g., collected stock market data or network packet data. It is intended that the current invention has application to all of these kinds of data for choosing an appropriate training set.

According to the present invention, available historic data is maintained in the form of snapshots of sensor or other parametric data, each snapshot having the same number of parametric data values. The parametric values are arranged as a vector, with corresponding parametric values as elements in the same position in each vector from snapshot to snapshot. That is, for example, the first element in all historic available data arranged as snapshot vectors may be designated to be the temperature reading from sensor A, the second element may be designated to always be the pressure reading from sensor B, etc.

For each sensor for which it is desirable to include more collected snapshots in the resulting training set to provide better characterization of the dynamics of the sensor in dynamic regimes, the range

of values of that sensor is determined over the set of all the collected snapshots. Alternatively, a range can be an expected range, based on knowledge of the application environment. This range is then divided into intervals. Preferably, in order to best capture the dynamic regions of the 5 range, the intervals are equal intervals, however they can be based on other criteria as well. The set of all collected snapshots is then reviewed, examining the value in each snapshot of that particular sensor, and one is selected for each interval, for inclusion in the training set. The value that falls closest to the upper or lower edge of the interval, or the middle of the 10 interval, can be used to determine which snapshot will be selected for each interval. The size of the interval should be selected in view of the total number of snapshots in the collection, as well as the desired amount of augmentation of a training set already populated by the Min-Max method. A preexisting training set derived from Min-Max training does not need to 15 be formed in order to use the present invention: The entire training set can be selected by means of this invention. In fact, when the intervals are set up such that they include the least and greatest bound of the data for a given sensor in the collected set, the Min-Max snapshots are thereby naturally included as a result.

20 The method of selecting snapshots for inclusion in the training set can be better understood by visually ordering the values of the sensor from the snapshots in increasing amount, and graphically indicating which are selected using equally spaced intervals.

25 After arranging the vectors in parameter value ascending or descending order, only the scalar value of the parameter of interest in that iteration is used. Fig. 5 graphically depicts a sequence of values as a bar chart ordered with increasing parameter height. In this example, the Y-axis 510 represents the magnitude of the particular parameter, e.g., temperature, 30 pressure or the like. The X-axis 520 is merely a sequential list of the vectors, i.e., sequence number. In the most general sense, historic snapshots are selected for inclusion in a training set according to the present invention, by iterating a selection procedure for each parameter or vector element. For each parameter, the entire available data set is arranged in an ascending or descending order of that particular parameter's magnitude.

Then, the magnitude along the Y-axis is then divided up
equally into bins, each represented in the present example by equally spaced
horizontal lines, e.g. 530, that cross the entire chart. Additional vectors 531,
532, 533, 534, 535, 536, 537, 538, 539 and 544 are selected for inclusion in the
5 acceptable state matrix D, if the scalar value in that vector for the particular
parameter of interest is closest to one of the bin boundaries 530. For each bin
boundary line 530, only one vector can be chosen. In the most general sense,
the scalar of interest (i.e. the parameter value) can be more than or less than
the bin boundary value, so long as the absolute value of the difference is the
10 smallest amount among all the scalars of interest.

15 In this example, those vectors that are selected for inclusion
according to this invention are shown in solid black. The bins intervals
along the Y-axis between lines are all equal, although this is not necessarily a
requirement.

15 Advantageously, using this method, the training set includes
more vectors where the sensor value of the sensor in question comes from
the range over which fewer vectors span a large y-axis regime. For example,
over the total range of values 540 for the sensor in the collected set, one
vector 544 is selected over the set of vectors 547 that have sensor values
20 falling around that value. Similarly, over the set of vectors 550, only one
vector is selected for inclusion. However, in the value range 553, only a
minimal set 560 of vectors populated the range, and most were selected for
inclusion, thereby better representing that range.

25 This binning procedure is carried out for each parameter or
vector element, selecting vectors for each. Then, the selected sets of vectors
for each parameter are combined, any duplicates are eliminated, and the
resultant set of vectors, or snapshots, is the preferred training set. This
preferred training set may be used in the empirical modeling described in
the aforementioned patents, or in any similar such system monitoring to
30 achieve enhanced monitoring sensitivity for any process, machine or system.

A bin interval may be along the y-axis chosen for each
parameter. Some system parameters may be known as dominant drivers in
the dynamics of the underlying system, process or machine, and so, it may
be preferable to emphasize those while deemphasizing others, including

selecting more vectors for those dominant drivers. In this case, the bin interval for dominant drivers may be of a much finer resolution than for other parameters, resulting in more vectors being selected, finely covering the dynamic range for dominant drivers, than for other more coarsely 5 monitored system parameters.

Turning to FIG. 6A, a block diagram of an exemplary laboratory workbench arrangement 600 is shown for gathering system, process or machine behavior data for distillation. In this example, the monitored system is depicted as a machine prototype 602 and may be, for 10 example, a combustion engine, an electric motor, a pump, a compressor, a refrigerator, and so on. It is understood that, as further indicated hereinabove, the monitored system may be any machine, living system or system carrying out a process. In this example, the machine 602 is labeled a prototype, but importantly, it should generate sensor data that is 15 substantially the same as the actual parameter values expected in a production model of the machine, as would be measured by the same sensors. Of course, the training may be in situ wherein the prototype is a production model itself, and ideally, not different in any way from other production models. In addition when sufficient system data has already 20 been accumulated that previously accumulated data may be used as the training data source, the prototype machine being a virtual machine derived from the production machine contributing data to the accumulation.

The machine 602 may be connected to and controlled by a control system 604, generally comprising a microcontroller- or 25 microprocessor-based digital system with appropriate analog/digital and digital/analog inputs and outputs are known to those skilled in the art. Machine 602 is instrumented with sensors monitoring machine components or reactions thereto (e.g., chamber temperature or pressure) and providing resultant sensor values along outputs 606. During training, the machine 604 is operated through an expected range of operations, and data acquisition 30 system 608 records values of all sensors 606 with which machine 602 is instrumented. Additionally, control signals from control system 604 may also be recorded by data acquisition system 608, and may be used as "sensor signals" that correlate with the other sensor signals.

Data acquired by data acquisition system 608 can accordingly be processed using a computer module 610 for producing a distilled training set of data representing the operational ranges of machine 602, using the training method described herein.

5 The monitoring system described herein includes an empirical modeling engine and a statistical decision-making engine supported by a suite of software routines for data preconditioning, training, and post-decision reporting. This system is modular and can be applied separately depending on the requirements of the particular monitoring application.

10 Typically, process monitoring equipment employs sensors having some common characteristics. A set of sensor data is acquired as being representative of the normal or desired operation range of the system which is made available for training as described for Fig. 3. The sensors chosen for the model should be correlated, either linearly or nonlinearly. Generally,

15 multiple sensor inputs may be necessary, however, the described algorithms may apply to single sensor applications by using signal decomposition of the sensor signal into components which can be treated as multiple, correlated inputs for modeling and monitoring. The identification of small deviations in signals from normal operation is provided as indicative of the

20 status of the sensor's associated physical parameters.

Thus, an evaluation system 650 with an on-board processor is shown in FIG. 6B, wherein a system, machine or process 652 is controlled by a control system 654 that is located on the machine. Machine 652 is instrumented with sensors for some of the physical or logical parameters of interest in controlling the machine, and the outputs for these sensors are shown as output conductors 656, which feed into the control system 654. These are also passed to a processor 658 located within or on the machine 652, disposed to execute a computing program for monitoring sensor signals and an optional computing program for generating a set 660 of virtual signals on the output conductors 656. The processor 658 is connected to a local memory 662, also on or in the machine 652, which stores data comprising the training set distilled according to the present invention to represent the expected operational states of the machine 652.

30 Advantageously, memory 662 can also store programs for execution by the

processor 658. Virtual signals 660, if included, previously generated by the processor 652 are provided to the control system 654, in lieu of genuine sensor values. Generation of virtual sensor estimates using the improved similarity operator of the present invention can be more fully understood 5 with reference to copending patent application no. 09/718,592 of Wegerich, filed November 22, 2000, and entitled "Inferential Signal Generation for Instrumented Equipment and Process." Virtual signals may be generated as a cost saving measure or 658 for unmonitorable physical or logical machine parameters.

10 Processor 658 can also be a part of the control system 654, and in fact can be the processor on which the control system routines are executed, in the event the control system 654 is a digital computer control system. Ideally, the processor 658 and memory 662 are powered by the same power source as the control system 654. However, under certain 15 circumstances, it may also be preferable to provide for a processor 658 and memory 662 independent from the processor and/or memory of the control system 654, in order to provide virtual signals 660 in a timely fashion, as though they were truly instrumented parameters. As an example, it may be necessary that processor 658 must operate at a higher clock speed than the 20 control system processor.

FIG. 7 shows an example 700 wherein a process 702 is instrumented with sensors having output leads 704. These leads 704 provide sensor signals to a control system 706 that controls the process 700. These signals 704 are also provided to a remote communications link 708, 25 which is disposed to communicate digital prior signal values to a second remote communications link 710, located at a physically remote place. A processor 712 may be included, which may act as a software controlled computer using the sensor signals received by link 710, optionally, to monitor the process 702 for sensor failures, process upsets or deviations 30 from optimal operation and optionally generate virtual sensor signals indicative of an inferred physical parameter of process 702. A memory 714 is provided to store training set data representative of the expected operational behavior of the process 702, selected according to the present invention.

Furthermore, a display 716 may be provided at the remote location for displaying data descriptive of the process 702, i.e. sensor signals 704 and any virtual signals derived therefrom or both. The virtual signals generated by processor 712 can also be transmitted from link 710 back to link 708 and input over leads 718 to control system 706 for advantageous control of the process. Data from original sensor signals and/or virtual sensor signals can also be transmitted to a third remote communications link 720, located at yet a third distant place, for display on display 722, thereby providing valuable information concerning the process to interested parties located at neither the physical process site nor at the site where optional virtual signals are computed and the system monitoring is processed.

The remote communications links can be selected from a variety of techniques known in the art, including internet protocol based packet communication over the public telecommunications infrastructure, direct point-to-point leased-line communications, wireless or satellite. More specifically, remote links 708, 712 and 720 may be internet-enabled servers with application software for accumulating, queuing and transmitting data as messages, and queues for receiving and reconstituting data arriving as messages. Alternatively, communications can be synchronous (meaning in contrast to asynchronous, message-based communications) over a wireless link.

Another embodiment of the invention is shown in FIG. 8, wherein is shown the cumulative distribution function 810 for the same data as in FIG. 5. The cumulative distribution function provides a measure of probability that a randomly selected vector from the vector population that comprises the collected data will have a sensor value for the sensor in question that is less than or equal to the x-axis value. Consequently, the curve reaches 1, or 100%, at the extreme right and begins at zero at the extreme left of the x-axis. The x-axis is scaled to the range of data observed for the sensor over the collected data. A steep section of curve 810 indicates a large number of vectors that have a sensor value x around the steep section, while a flat region indicates a region of the x-axis sparsely populated in the collected data set. Lines 820 indicate equally spaced intervals that can be used to select vectors for inclusion in the training set. Wherever a line

820 intersects the curve 810, a vector must be chosen to add to the training set. Typically, for each such point, a subset of the collected snapshots will comprise those that make up the CDF, that is, are those that have a sensor value less than or equal to the value x along the x-axis. From that set, the 5 vector having the highest value for that sensor is chosen.

The effect of this embodiment of the invention is to select more vectors from sub-ranges of the range observed for a sensor, where the data is heavily populated, such as at sub-range 830. Sub-ranges like 840 where data is sparse, are also sparsely represented in the final training set. This effect is 10 overall opposite of the effect of that shown in FIG. 5, and is useful in the event that the model needs finer gradation of estimates for certain ranges where the data is densely populated.

FIG. 9 is a flow diagram 900 for the preferred embodiment of the present invention. Historical data is collected in step 902 as described 15 hereinabove. Optionally, in step 904 dominant drivers may be identified. Then, in step 906 a parameter is selected from parameters included for collected sensor data. In step 908 historical data is ordered, sorting snapshots according to the selected parameter, e.g., in ascending order for that parameter. Next, in step 910 the vector space defined by the selected 20 parameter is binned, i.e., bins are defined for the selected parameter. If additional snapshots are to be included only for dominant drivers, then for nondominant drivers the minimum and maximum of that particular parameter are provided as bin intervals in step 910. In step 912 vectors are selected for inclusion in the training set, for example selecting those vectors 25 where the parameter value most closely approaches a bin value for inclusion in the training set. In step 914 a check is made to determine if any parameters remain unrepresented, i.e., have not had at least a min and max vector identified therefor. If any parameters remain, then returning to step 906 one of the remaining parameters is selected. Otherwise, vector selection 30 is complete and in step 916 redundant vectors that may have been included are eliminated from the selected vectors. Next, in step 918 the selected vectors are stored as the training set. Finally, in step 920 training is complete.

Each selected vector is included in step 918 in its entirety in the training set regardless of other parameter values in that vector that might not approach any corresponding bin interval. Thus, if a vector is chosen in step 912 because of a particular parameter value being closest to a bin 5 interval, the entire vector becomes part of the training set. With a sufficiently fine bin interval, the minima and maxima of all parameters are included in the final training set in step 918, just as would be achieved by applying the Min-Max method described above.

Numerous alternate tests may be substituted for use in 10 selecting vectors in each bin in step 912 for inclusion in the training set. For example, those vectors wherein the parameter value that is closest to, but does not exceed, the bin value, or, conversely, vectors with parameter values that are closest to, but in no case less than the bin value may be included. In another example, the bin intervals may be chosen in a way that varies, as 15 opposed to being constant, such as selecting a mathematical function that describes the bin intervals, e.g., a logarithmic function, or a geometric correspondence. Further, any of a wide variety of mathematical dependencies may be chosen for use in this invention as might readily be understood by those skilled in the art.

In yet another example, instead of using magnitudinal bin 20 intervals to select training set vectors, a periodic or an aperiodic bin interval along the x-axis can be chosen, those vectors falling closest to the bin intervals being included in the training set. It should be noted that because the x-axis is an unitless enumerated axis indicating snapshot sequence 25 number and with no or at best a very attenuated relationship to the vector, for purposes of ordering, the vectors may be ordered according to increasing or decreasing magnitude of the particular parameter value is analogous to selecting according to bin value, when selecting every nth vector in the same order, where n is some preselected number, especially for well behaved 30 parameters.

It should be noted that although selection can be done graphically as described hereinabove, the present invention does not require that a chart be created and displayed; rather, FIGS. 5 and 8 are provided for example only as a way of describing visually what can be carried out

computationally. By way of example, all the vectors in the historic available data can be maintained in a database and sorted according to parameter value. Then, bins may be calculated for the entire parameter range or for the entire parameter magnitude. Then, stepping through the database records 5 in order, the parameter value for each vector is compared against a current bin value. When the closest parameter value is found, that vector is marked for inclusion in the training set. This is repeated until a vector is identified for each bin.

As indicated hereinabove, with reference to Figs. 6A-B and 7, 10 the present invention can be carried out on a computer with a memory and processor executing software to perform the necessary computations to generate the memory-stored file of resulting vectors in the final training set. The system can also include a training interface receiving real-time parameter data from sensors or from a distributed control system or the like. 15 Therefore, the training set can be developed coincident with viewing a real-time signal feed from one or more sensors for a particular time period. Alternatively, the computer can sift through historic data stored in a file providing a training set data file therefrom.

In another embodiment, a microprocessor coupled with 20 sufficient memory to store the historic data, either on-board or off-board, can be controlled to store the resulting training set locally for use with monitoring activities such as those described hereinabove.

It should be appreciated that a wide range of changes and 25 modifications may be made to the embodiments of the invention as described herein. Thus, it is intended that the foregoing detailed description be regarded as illustrative rather than limiting and that the following claims, including all equivalents, are intended to define the scope of the invention.